REMARKS

Claims 1-9, 11-17, 19-38, 40-44, 46-50, 52-60 are now pending in the application. Claims 10, 18, 39, 45, 51, and 57 have been canceled without prejudice or disclaimer. Claims 1, 3, 5, 11, 12, 13, 16, 17, 19, 34, 40, 46, and 52 have been amended, and new claims 59-60 added, without introduction of new matter. Favorable reconsideration is respectfully requested in view of the above amendments and the following remarks.

Claims 1-11, 17-22, 32, 34, 38, 39, 40, 44-46, 50, 51, 52, 56, and 57 stand rejected under 35 U.S.C. §§102(a) and 102(b) as allegedly being anticipated by Childers et al. (US Patent 5,986,913). This rejection is respectfully traversed.

Independent claims 10 and 18, as well as dependent claims 39, 45, 51, and 57 have been canceled, thereby rendering the rejection of these claims moot. Claims 11, 34, 40, 46, and 52, which had depended from claim 10, have each been amended to now ultimately depend from independent claim 1 via new dependent claim 60. Formerly independent claim 17 has also been amended so that it now depends from claim 1.

Before delving into the details of the Office Action, it is believed that a discussion of the invention in its broadest sense will yield a better understanding of various aspects pertinent to patentability.

A processor in accordance with the invention is formed of a plurality of processing elements (PEs) adapted to receive an incoming stream of data packets, which may be of unpredictable length. There is no requirement for an incoming packet stream in which every packet is of the same fixed size. In the case of a small packet, it can be read into a single processing element along with its header. Where an incoming packet is large, it is "chopped up" into smaller fragments or "chunks" of substantially equal size and spread over as many PEs as are necessary to store it and process it. This is spelled out in the description and is specified in original claim 4. The product of packet size and packet rate is invariably constant (see the specification at page 7, lines 19-21), so there is no possibility of the incoming stream having large packets and a high packet rate at the same time.

Thus, each chunk is sent to a respective PE. If an incoming large data packet is not an integral multiple of the size of a chunk, there will be a residual chunk that is not the same size as the rest. This residual chunk will nevertheless be sent to a PE along with the others. It could be padded out with zeroes to make it the same size as all the other chunks but its length could simply be indicated in some way, such as in a header for that chunk.

Various passages throughout the specification support these various characteristics, such as page 12, line 16; page 13, lines 6-9; page 22, lines 19-21; page 23, lines 26-32; page 24; page 51; and page 59. Claim 1 of course relies on these characteristics. In addition, dependent claim 4 specifically relates to the aspect that the chunks are of a fixed size and are consequently equal in size. Significantly, page 59, lines 15-20 confirm that the processor is matched to a constant line bandwidth rather than to any particular number of packets. The significance of this will become apparent when discussing the bandwidth restriction in claim 1 and the cited prior art.

The wording of original claim 1, when referring to distributing data packets "in whole or in part", was intended to encompass two important aspects. First, the concept that packets are allocated in their entirety to the PEs. In numerous prior art systems, the header is stripped from a packet and sent to a processor while the information-containing part of the packet is stored separately. The problem with this prior art structure is that the processed header has to be matched up with the stored information-containing part of the packet after processing. Second, small packets are distributed "whole" (i.e., an entire (small) packet is sent to a single PE), whereas a large packet is distributed "in parts" over a number of PEs so that each of those PEs stores only a part (a chunk) of the packet.

In order to make this distinction even more clear, the revised wording of claim 1 now states that "for any given data processing bandwidth of the processing elements the input device is operable to distribute whole data packets across one or more processing elements". Thus, small packets can be sent whole to one PE and large packets can be subdivided and distributed across a number of PEs, as described in the passages of the description mentioned above and as explicitly claimed in original dependent claims 3, 4, and 5, for example. The reference to the bandwidth aspect of the invention has also been revised in claim 1 to even more clearly define, in keeping with the description, that for any given data processing bandwidth of the PEs, whole packets or chunks are distributed to one or more PEs.

Turning now to specific aspects of the Office Action, the Office has cited Childers et al. against a number of claims, including original claim 1. The Childers et al. patent discloses a high-speed sense amplifier/ memory configuration resistant to voltage spikes. The practical embodiment is a serial video processor in which lines of video are encoded to create packets representing respective pixels of a video line. The processor is a SIMD/RISC device consisting of a one-dimensional array of 1-bit PEs shown in Figure 1 as a parallel processor.

Figure 2 represents a block diagram of a single PE of the processor. There are as many PEs as there are pixels in the video line. Data for each pixel is handled by a respective PE. The packets are all of the same size. There is considerable discussion as to the number of bits for the PEs and the fact that there are up to 1024 words x 40-bits of input data but it seems inevitable that the allocation of data to the PEs is based purely on numbers of pixels and numbers of PEs.

In Childers et al., the processor has clearly been engineered to match the number of pixels with the number of PEs on a one-to-one basis. In contrast, a processor architecture in accordance with claim 1 operates in ignorance of the size of packets in the incoming data stream but automatically subdivides a larger packet into fragments that are then allocated to as many PEs as are necessary to store it for a given bandwidth. This allocation is based on data bandwidth rather than pixel number. It is therefore respectfully asserted that Childers et al. does not anticipate the features defined by claim 1.

The remaining claims 2-9, 11, 17, 19-22, 32, 34, 38, 40, 44, 46, 50, 52, and 56 all depend from claim 1, and are therefore patentably distinguishable over Childers et al. for at least the same reasons as set forth above, as well as for the additional features that they define which are also not disclosed in the Childers et al. patent.

In view of the above, it is respectfully requested that the rejection of claims 1, 2-9, 11, 17, 19-22, 32, 34, 38, 40, 44, 46, 50, 52, and 56 under 35 U.S.C. §§102(a) and 102(b) be withdrawn.

Claims 12, 13, 35, 41, 47, and 53 stand rejected under 35 U.S.C. §§102(a) and 102(b) as allegedly being anticipated by Gove et al. (US Patent 5,371,896). This rejection is respectfully traversed.

Formerly independent claim 12 has been amended so that it now depends from claim 1, and thereby incorporates all of the features of that base claim. Claim 12 is not anticipated by Gove et al. at least because Gove et al. fail to disclose or suggest the various features discussed above with respect to claim 1. Moreover, Gove et al. disclose a processor switchable between SIMD and MIMD using a cross-bar switch interconnecting various PEs and memory to change the combination of distributed/shared memory. There are several parallel processors interconnected with the memory banks. When there is contention (i.e., simultaneous accesses to RAM by any two system devices) a "SIMD pause" signal is routed to pause *all* PEs (see Gove et al. Fig. 30). The pause is therefore not dependent on data, as

required by the Applicants' claim 12. The standby mode required of Applicants' claim is therefore not the same as the global "pause" in Gove et al.

Claims 13, 35, 41, 47, and 53 depend from claim 12, and are therefore patentably distinguishable over the Gove et al. patent for at least the same reasons as those set forth above.

In view of the foregoing, it is respectfully asserted that claims 12, 13, 25, 41, 47, and 53 are patentable over the prior art of record. Accordingly, it is respectfully requested that the rejection of these claims under 35 U.S.C. §§ 102(a) and 102(b) be withdrawn.

Claims 23-31 and 58 stand rejected under 35 U.S.C. §§ 102(a) and (b) as allegedly being anticipated by Horst (US Patent 5,404,550). This rejection is respectfully traversed.

A concept encapsulated in independent claim 23 involves an I/O system for transferring data to and from PEs in a SIMD array such that data packets of different sizes are sent to respective ones of the PEs in the array. In other words, the PEs are large enough to accept whole packets.

The Office relies on Horst as being relevant to the novelty of claim 23. This reliance is unfounded for a number of reasons. For one thing, Horst discloses inter-processor networks as opposed to I/O systems. Horst is not SIMD. Instead, it operates by sending transmission packets between PEs, each such packet carrying the state of execution from the previous PE. In the first embodiment, the PEs are interconnected unidirectionally.

Horst has a mechanism for routing message blocks between PEs. If there is no direct path, then the messages are routed through intervening PEs to their destination. Neither of these is explicitly for I/O, although presumably an I/O port could be the final destination of a message. In both cases the message blocks are of a fixed size and format (address, control information, data, etc.). In contrast, the I/O mechanism of claim 23 is shared by (connected to) all PEs. Also, the chunks distributed to each PE just contain a "slice" from the incoming data packet, so there is no need for routing information and the like.

Claims 24-31 and 58 variously depend from claim 23, and therefore incorporate all of the features of that base claim. Accordingly, for the foregoing reasons claims 23-31 and 58 are believed to be patentably distinguishable over the Horst patent. It is therefore respectfully requested that the rejection of these claims under 35 U.S.C. §§ 102(a) and 102(b) be withdrawn.

Claims 14, 15, 36, 37, 54, and 55 stand rejected under 35 U.S.C. §103(a) as allegedly being unpatentable over Brown (US Patent 5,872,993) in view of Childers et al. This rejection is respectfully traversed.

In order to establish a *prima facie* case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, and not based on applicant's disclosure. *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991).

The Office has failed to make out a *prima facie* case of obviousness for a number of reasons. To begin with, the necessary suggestion or motivation for making the combination described in the Office Action is lacking in the prior art. More particularly, in Figure 3, Brown discloses a hardware accelerator that communicates with a processor via an interface/"Bus Master" 320 which "processes" data streams. Memories in the DSP are used as buffers whereas the accelerator uses memory 350B. As acknowledged on page 5 of the Office Action, the device is not described as using a SIMD architecture.

The Office attempts to make up for the deficiencies of Brown by arguing that it would have been obvious to apply the Hardware accelerator concept in Brown to the SIMD architecture taught in Childers et al.

Applicants respectfully disagree. In a SIMD processor, it only makes sense to pass data from a parallel processor to a hardware accelerator if that accelerator is also parallel, otherwise there is a "serialization problem". This is something to be avoided in parallel processors because after data items from each PE have been sent to the accelerator, it is necessary to wait for that item of data to be processed and returned before the next data item from another PE can be sent. The whole point of parallel processing would then be lost. Thus, the prior art fails to provide the necessary suggestion or motivation for making the combination now proposed by the Office.

Furthermore, it is not at all apparent that one of ordinary skill in the art would have had a reasonable expectation of success in making the Office's proposed combination. The

obstacles to combining Childers and Brown, which have been overcome by the invention as set out in claims 14 and 15, can be summarized as follows:

First, a more complex accelerator is required which, although it may receive data sequentially, will have to be parallel internally (typically pipelined) so it can process several of the requests at once. At any point in time, there may be several data items queued up waiting to be processed by the accelerator, several being processed and several being returned to the PEs.

Second, in order to avoid inefficiencies for the SIMD processor, it must be able to do something else while this is happening. This requires (a) multi-threading and (b) an I/O mechanism that allows the data to be returned to the PE memory while the PEs are otherwise involved with something else, that is, effectively without their direct intervention. This is significantly more complex than a sequential processor where the data could just be returned to shared memory. On the contrary, the memory of PEs in a SIMD processor is typically "private" (to each PE) rather than shared.

Third, communication between the SIMD processor and the accelerator needs a mechanism to "'tag" each data item with the identity of the originating PE so it can be returned to the correct place. In this connection it is important to note that because the SIMD processor will not return to processing the data returned by the accelerator until all the data has been returned (or there is a timeout) there is no need to re-order the data if it is returned out of order as long as it is returned to the correct PE.

It should therefore be apparent that, far from being a "simple" operation to combine the teachings of the two cited documents, it is actually a complex situation that raises as many problems as it "solves". It leaves unanswered questions as to how to design mechanisms by which the SIMD processor can interact with the accelerator in an effective and efficient way. Claims 14 and 15 are concerned with two particular ways of achieving this. In claim 14 the accelerator is operable to return processing results (from the accelerator) when those processing results are available. In claim 15, the accelerator is operable to return processing results to respective processing elements in the order in which processing requests were received by the accelerator.

Claims 16, 33, 36, 37, 42, 43, 48, 49, 54, and 55 variously depend from claims 14 and 15, and consequently incorporate the same features as those defined by their respective base claims. It is therefore respectfully submitted that claims 14, 15, 16, 33, 36, 37, 42, 43, 48, 49,

54, and 55 are patentably distinguishable over the Brown and Childers et al. patents regardless of whether these documents are considered individually or in any combination. Accordingly, it is respectfully requested that the rejection of these claims under 35 U.S.C. § 103(a) be withdrawn.

New claims 59 and 60 have been added without introduction of new matter. These claims each depend from claim 9, and are therefore patentable for at least the same reasons as those set forth above.

The application is believed to be in condition for allowance. Prompt notice of same is respectfully requested.

Respectfully submitted,
Potomac Patent Group PLLC

Date: April 18, 2005

By: 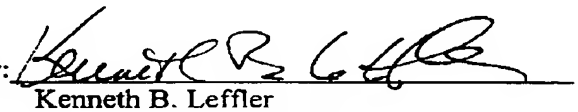
Kenneth B. Leffler
Registration No. 36,075

P.O. Box 855
McLean, Virginia 22101-0855
703-718-8884

I hereby certify that this correspondence is being sent by facsimile transmission to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 to the following facsimile number:

Facsimile Number: 703 872 9306
Date of Transmission: April 18, 2005


Kenneth B. Leffler